



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/650,800	08/30/2000	Paul S. Neuman	RA 5290(33012/289/101)	1186
27516	7590	05/10/2005	EXAMINER	
UNISYS CORPORATION MS 4773 PO BOX 64942 ST. PAUL, MN 55164-0942			VITAL, PIERRE M	
			ART UNIT	PAPER NUMBER
			2188	

DATE MAILED: 05/10/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

09/650,800

Applicant(s)

NEUMAN, PAUL S.

Examiner

Pierre M. Vital

Art Unit

2188

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 21 March 2005.
- 2a) ☒ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-25 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-25 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 30 August 2000 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☐ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
Paper No(s)/Mail Date _____
- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____

DETAILED ACTION

Continued Examination Under 37 CFR 1.114

1. A request for continued examination under 37 CFR 1.114, including the fee set forth in 37 CFR 1.17(e), was filed in this application after final rejection. Since this application is eligible for continued examination under 37 CFR 1.114, and the fee set forth in 37 CFR 1.17(e) has been timely paid, the finality of the previous Office action has been withdrawn pursuant to 37 CFR 1.114. Applicant's submission filed on March 21, 2005 has been entered.

Response to Amendment

2. This Office Action is in response to applicant's communication filed March 21, 2005 in response to PTO Office Action mailed February 15, 2005. The Applicant's remarks and amendments to the claims and/or the specification were considered with the results that follow.

3. Claims 1-25 have been presented for examination in this application. In response to the last Office Action, claims 5-6, 9, 11, 14-15 and 19 have been amended. No claims have been canceled or added. As a result, claims 1-25 remain pending in this application.

Response to Arguments

4. Applicant's arguments, see Remarks, filed March 21, 2005, with respect to the rejection(s) of claim(s) 1-4, 6, 11 and 21 under 35 USC 103(a) have been fully considered and are persuasive. Therefore, the rejection has been withdrawn.

However, upon further consideration, a new ground(s) of rejection is made in view of Applicant Admitted Prior Art (hereinafter "AAPA").

5. Applicant's arguments filed August 19, 2004 have been fully considered but they are not persuasive. As to the remarks, applicant asserted that:

(a) Lai does not teach or suggest the "first checking" and "second checking" steps.

Examiner respectfully traverses applicant's arguments for the following reasons. Examiner would like to point out that the fact that the "checking steps" in Lai is redundant because of the inclusion policy, it further reinforces that the checking is done in both the first and second level caches because a hit in L1 cache will result in a hit in L2 cache.

(b) There is no showing that Lai "records the location of data corresponding to read memory request within a level one cache memory" following a level one cache memory miss.

Examiner respectfully traverses applicant's arguments for the following reasons. Examiner would like to point out that the victim line recorded in Lai is a least recently

used (LRU) line chosen when a primary cache (i.e., L1 cache) read miss occurs and stored in a victim buffer (col. 5, lines 22-30). Thus, by displacing the LRU line and storing the line in a buffer, it can be clearly seen that the address or location of the line is recorded as claimed by applicant. Furthermore, Lai discloses storing data in primary (i.e., L1) cache when a miss to L1 cache and L2 cache occurs (col. 6, line 65 – col. 7, line 11).

(c) The Examiner used mutually exclusive embodiments as taught by Lai.

Even though Lai discloses two different cache configurations, they are not mutually exclusive as asserted by Applicant. Note that Lai discusses using a partially inclusive cache as well (col. 6, lines 9-41). Because the prior art structure is capable of performing the intended use, then it meets the preferred embodiment of Lai.

(d) There is no showing of the generation of a parity error by a level two cache memory in Hazawa.

Note that the level 0 cache of Hazawa is equivalent to the level 1 cache of the present invention and the level 1 cache of Hazawa is equivalent to the level 2 cache of the present invention and so forth. Also note that Hazawa discloses an error has occurred at level 1 cache (col. 3, lines 45-57). Invalidating can only be done at corresponding locations of the L1 and L2 caches since parity check is often used to check the accuracy with which each byte is stored. As such, Hazawa teaches the generation of a parity error by a level two cache memory.

(e) Lai describes SNOOPing of a remote processor memory request rather than a cache first and second memory request.

In response to applicant's argument that the references fail to show certain features of applicant's invention, it is noted that the features upon which applicant relies (i.e., SNOOPing of a remote processor memory request rather than a cache first and second memory request) are not recited in the rejected claim(s). Although the claims are interpreted in light of the specification, limitations from the specification are not read into the claims. See *In re Van Geuns*, 988 F.2d 1181, 26 USPQ2d 1057 (Fed. Cir. 1993). Claim 16 does not at all disclose snooping of a cache or of a processor.

Claim Rejections - 35 USC § 112

6. The following is a quotation of the first paragraph of 35 U.S.C. 112:

The specification shall contain a written description of the invention, and of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same and shall set forth the best mode contemplated by the inventor of carrying out his invention.

7. Claims 11-15 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the enablement requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to enable one skilled in the art to which it pertains, or with which it is most nearly connected, to make and/or use the invention. Claim 11, as amended, recites the limitation "a semi-inclusive level two cache memory" in lines 3-4. However, this limitation is not disclosed in the page and line numbers the applicant provided to support this amendment.

Claim Rejections - 35 USC § 102

8. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

9. Claims 16 and 19 are rejected under 35 U.S.C. 102(b) as being anticipated by Lai (US5,564,035).

As per claim 16, Lai discloses an apparatus comprising:

executing means for executing program instructions [*processor 200, 210; Fig. 3*]; level one caching means responsively coupled to said executing means for level one caching data [*L1 cache 202; Fig. 3*]; accessing means responsively coupled to said executing means and said level one caching means for accessing a data element if said executing means requires accessing of said data element [*system bus 208; Fig. 3*]; level two caching means responsively coupled to said requesting means for level two caching data [*L2 cache 203; Fig. 3*]; and first invalidating means responsively coupled to said level one caching means for invalidating said data element if said data element is a write data element located within said level two caching means and within said level one caching means [*hit in L2 cache and hit in L1 cache, corresponding cache line can be invalidated; col. 2, lines 52-57*].

As per claim 19, Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency [col. 6, line 45 – col. 7, line 11].

Claim Rejections - 35 USC § 103

10. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

11. Claim 1 is rejected under 35 U.S.C. 103(a) as being unpatentable over Applicant Admitted Prior Art (hereinafter "AAPA") and Lai (US5,564,035).

As per claim 1, drafting a claim in Jepson format (i.e., the format described in 37 CFR 1.75(e); see MPEP § 608.01(m)) is taken as an implied admission that the subject matter of the preamble is the prior art work of another. In re Fout, 675 F.2d 297, 301, 213 USPQ 532, 534 (CCPA 1982) (holding preamble of Jepson-type claim to be admitted prior art where applicant's specification credited another as the inventor of the subject matter of the preamble). However, this implication may be overcome where applicant gives another credible reason for drafting the claim in Jepson format. In re

Art Unit: 2188

Ehrreich, 590 F.2d 902, 909-910, 200 USPQ 504, 510 (CCPA 1979) (holding preamble not to be admitted prior art where applicant explained that the Jepson format was used to avoid a double patenting rejection in a co-pending application and the examiner cited no art showing the subject matter of the preamble). Moreover, where the preamble of a Jepson claim describes applicant's own work, such may not be used against the claims. Reading & Bates Construction Co. v. Baker Energy Resources Corp., 748 F.2d 645, 650, 223 USPQ 1168, 1172 (Fed. Cir. 1984); Ehrreich, 590 F.2d at 909-910, 200 USPQ at 510.

However, AAPA does not specifically teach a first logic which invalidates a corresponding level one cache memory location in response to either a non-local write as recited in the claim.

Lai discloses invalidating a corresponding level one cache memory location in response to a non-local write [*when second processor performs write to memory, hit in L1 cache, corresponding cache line can be invalidated*; col. 2, lines 41-57] to maintain multiprocessor coherency. Since the technology for implementing invalidating a corresponding level one cache memory location in response to a non-local write was well known and since invalidating a corresponding level one cache memory location in response to a non-local write maintains multiprocessor coherency, an artisan would have been motivated to implement invalidating a corresponding level one cache memory location in response to a non-local write in the system of Mayfield and Lynch. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify

the system of Mayfield and Lynch to include invalidating a corresponding level one cache memory location in response to a non-local write because it was well known to maintain multiprocessor coherency as taught by Lai.

12. Claim 2-4 are rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Lai (US5,564,035) and Lynch et al. (US6,061,766).

As per claim 2, the combination of AAPA and Lai does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in on-chip cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit*; Fig. 4; col. 4, lines 19-30; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34). Note that mode 3 requests without ownership is described as meaning that there is no data within the level one cache memory to invalidate.

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an

artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of AAPA and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Lynch before him at the time the invention was made, to modify the system of AAPA and Lai to include implementing inhibiting a logic from invalidating mode 3 requests without ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 3, the combination of AAPA and Lai does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of AAPA and Lai. Thus, it would have been obvious to one of

Art Unit: 2188

ordinary skill in the art, having the teachings of AAPA and Lai and Lynch before him at the time the invention was made, to modify the system of AAPA and Lai to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 4, Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency (col. 6, line 45 – col. 7, line 11).

Since the technology for implementing recording location of data in response to a level one cache read miss and a level two cache memory read miss was well known and since recording location of data in response to a level one cache read miss and a level two cache memory read miss maintains multiprocessor coherency, an artisan would have been motivated to implement recording location of data in response to a level one cache read miss and a level two cache memory read miss in the system of AAPA. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of AAPA to include recording location of data in response to a level one cache read miss and a level two cache memory read miss because it was well known to maintain multiprocessor coherency as taught by Lai.

Art Unit: 2188

13. Claim 5 is rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Lynch et al. (US6,061,766) and Lai (US5,564,035) and Hazawa (US4,891,809).

As per claim 5, the combination of AAPA and Lai and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, AAPA and Lai and Lynch do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48]. Note that parity check is often used to check the accuracy with which each byte is stored as is well known by those of ordinary skill in the art.

It would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lynch and Hazawa before him at the time the invention was made, to modify the system of AAPA and Lynch to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

14. Claims 6 and 9 are rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Lai (US5,564,035) and Duncan et al (US 6,128,711).

Claim 6 is rejected using the rationale as for the rejection of claim 1 above. However, AAPA and Lai do not specifically teach a level two cache memory dedicated to the processor as recited in the claim.

Duncan discloses a processor chip 40 including a secondary cache 48 is included on-chip (Fig. 2; col. 6, lines 1323).

Since the technology for implementing a level-two cache on-chip was well known as evidenced by Duncan, an artisan would have been motivated to implement this feature in the system of AAPA and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Duncan before him at the time the invention was made, to modify the system of AAPA and Lai to include a level two cache memory dedicated to a processor because a level two cache included on chip provides data to the respective instruction and data stores to reduce the time required to obtain data from external memory (col. 6, lines 20-23) as taught by Duncan.

As per claim 9, AAPA does not specifically teach fourth logic which records location of data in response to a level one cache read miss and a level two cache memory read miss as recited in the claim.

Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency (col. 6, line 65 – col. 7, line 11). Since the technology for implementing recording in a

level one cache memory location of data in response to a level one cache read miss and a level two cache memory read miss was well known as evidenced by Lai, an artisan would have been motivated to implement this feature in the system of AAPA. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of AAPA to include recording location of data in response to a level one cache read miss and a level two cache memory read miss because it was well known to maintain multiprocessor coherency and to maintain cache coherency between the primary cache and main memory (col. 4, lines 7-10) as taught by Lai.

15. Claims 7 and 8 are rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Lai (US5,564,035) and Duncan et al (US 6,128,711) and Lynch et al. (US6,061,766).

As per claim 7, the combination of AAPA and Lai and Duncan does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit*; Fig. 4; col. 4, lines 19-30; *it is clearly obvious that any computer system*

uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system] to provide a snoop process for ensuring cache coherency (col. 2, lines 34). Note that mode 3 requests without ownership is described as meaning that there is no data within the level one cache memory to invalidate.

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of AAPA and Lai and Duncan. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Duncan and Lynch before him at the time the invention was made, to modify the system of AAPA and Lai and Duncan to include implementing inhibiting a logic from invalidating mode 3 requests without ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claim 8, the combination of AAPA and Lai and Duncan does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a*

combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of AAPA and Lai and Duncan. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Duncan and Lynch before him at the time the invention was made, to modify the system of AAPA and Lai and Duncan to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

16. Claim 10 is rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Lai (US 5,564,035) and Duncan et al (US 6,128,711) and Hazawa (US 4,891,809).

As per claim 10, the combination of AAPA and Lai and Duncan discloses the claimed invention as detailed above in the previous paragraphs. However, AAPA and

Lai and Duncan do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48].

It would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Duncan and Hazawa before him at the time the invention was made, to modify the system of AAPA and Lai and Duncan to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

17. Claims 11 and 14 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lai (US5,564,035) and Duncan et al (US 6,128,711).

As per claim 11, Lai discloses a method of maintaining validity of data within a level one cache memory of a processor responsively coupled to a level two cache memory which is responsively coupled to a system memory bus [*L1 cache 202, L2 cache 203, system bus 208*; Fig. 3] comprising: formulating a write request [*processor 210 performs a write to memory*; col. 2, lines 48-49]; first checking for a level one cache memory hit in

response to said write memory request [*hit in L1 cache 202*; col. 2, lines 48-52; col. 1, lines 60-62]; second checking a level two cache memory hit in response to a hit found by said first experiencing step [*L2 cache 203 detects a hit*; col. 2, lines 48-49; col. 2, lines 35-37]; and invalidating a portion of said level one cache memory corresponding to said write memory request in response to a hit found by said second checking step [*hit in L1 cache corresponding cache line can be invalidated*; col. 2, lines 52-57].

Duncan discloses a processor chip 40 including a secondary cache 48 is included on-chip (Fig. 2; col. 6, lines 1323).

Since the technology for implementing a level-two cache on-chip was well known as evidenced by Duncan, an artisan would have been motivated to implement this feature in the system of AAPA and Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Lai and Duncan before him at the time the invention was made, to modify the system of AAPA and Lai to include a level two cache memory dedicated to a processor because a level two cache included on chip provides data to the respective instruction and data stores to reduce the time required to obtain data from external memory (col. 6, lines 20-23) as taught by Duncan.

As per claim 14, Lai discloses recording location of data in response to a level one cache read miss and a level two cache memory read miss to maintain multiprocessor coherency (col. 6, line 65 – col. 7, line 11)

18. Claims 12-13 and 17-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lai (US5,564,035) and Lynch et al. (US6,061,766).

As per claims 12 and 17, Lai does not specifically teach a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership as recited in the claim.

Lynch discloses a second logic which inhibits said first logic from invalidating for mode 3 requests without ownership [*snoop requests checks for the presence of an object in cache; only requests for exclusive use which match cache tags are invalidated; there is no invalidate done when there is no hit*; Fig. 4; col. 4, lines 19-30; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing inhibiting a logic from invalidating mode 3 requests without ownership was well known and since inhibiting invalidating mode 3 requests without ownership provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement inhibiting invalidating mode 3 requests without ownership in the system of Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Lynch before him at the time the invention was made, to modify the system of Lai to include implementing inhibiting a logic from invalidating mode 3 requests without ownership because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

As per claims 13 and 18, the combination of Lai does not specifically teach third logic which invalidates said corresponding cache memory location in response to a SNOOP hit as recited in the claim.

Lynch discloses invalidating a corresponding cache memory location in response to a SNOOP hit [Fig. 4; col. 4, lines 24-26; *it is clearly obvious that any computer system uses a combination of logic to produce output based on the rules of logic it is designed to follow; clearly, the use of multiple logics is an inherent feature of any computer system*] to provide a snoop process for ensuring cache coherency (col. 2, lines 34).

Since the technology for implementing invalidating a corresponding cache memory location in response to a SNOOP hit was well known and since invalidating a corresponding cache memory location in response to a SNOOP hit provides a snoop process for ensuring cache coherency, an artisan would have been motivated to implement invalidating a corresponding cache memory location in response to a SNOOP hit in the system of Lai. Thus, it would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Lynch before him at the time the invention was made, to modify the system of Lai to include invalidating a corresponding cache memory location in response to a SNOOP hit because it was well known to provide a snoop process for ensuring cache coherency as taught by Lynch.

Art Unit: 2188

19. Claims 15 and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over Lai (US5,564,035) and Hazawa (US4,891,809).

As per claims 15 and 20, the combination of Lai and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, Lai and Lynch do not specifically teach invalidating a level one cache memory in response to a level two cache memory generating a parity error as recited in the claim.

Hazawa discloses invalidating a level one cache memory in response to a level two cache memory generating a parity error [col.3, lines 38-48].

It would have been obvious to one of ordinary skill in the art, having the teachings of Lai and Hazawa before him at the time the invention was made, to modify the system of Lai to include invalidating a level one cache memory in response to a level two cache memory generating a parity error because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa by the provision of a sequential verification logic circuit for generating error indicating signals in sequence within the cache memory unit [col. 1, lines 30-37] as taught y Hazawa.

20. Claim 21 is rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Hazawa (US4,891,809).

Claim 21 is rejected using the rationale as for the rejection of claim 1 above.

However, AAPA does not specifically teach a data element having a parity error stored in said level two cache memory and a facility responsively coupled to said level one cache memory and said level two cache memory which detects said parity error of said data element and invalidates a corresponding data element within said level two cache memory as recited in the claim.

Hazawa discloses invalidating data in a level two cache memory in response to a parity error of a data element to provide a cache memory with an error checking mode [col.3, lines 38-51; col. 3, lines 38-59].

It would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Hazawa before him at the time the invention was made, to modify the system of Mayfield to include invalidating data in a level two cache memory in response to a parity error of a data element because it was well known to provide a cache memory having a normal error checking mode as taught by Hazawa.

21. Claim 22 is rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Hazawa (US4,891,809) and Lynch et al (US6,061,766).

As per claim 22, the combination of AAPA and Hazawa discloses the claimed invention as detailed above in the previous paragraphs. However, AAPA and Hazawa do not specifically teach a level one cache memory comprising a level one instruction cache memory and a level one operand cache memory as recited in the claim.

Lynch discloses a level one cache comprising a data cache for storing data as it is passed back and forth from the execution units of the processor and an instruction cache holding instructions prior to execution by the processor's execution units (col. 3, lines 43-48).

It would have been obvious to one of ordinary skill in the art, having the teachings of AAPA and Hazawa and Lynch before him at the time the invention was made, to modify the system of AAPA and Hazawa to include a level one instruction cache memory and a level one operand cache memory because a level one data cache was well known for storing data as it is passed back and forth from the execution units of the processor and a level one instruction cache was well known for holding instructions prior to execution by the processor's execution units as taught by Lynch.

22. Claims 23-25 are rejected under 35 U.S.C. 103(a) as being unpatentable over AAPA and Hazawa (US4,891,809) and Lynch et al (US6,061,766) and Lai (US5,564,035).

As per claim 23, the combination of AAPA and Hazawa and Lynch discloses the claimed invention as detailed above in the previous paragraphs. However, AAPA and Hazawa and Lynch do not specifically teach invalidating a write data element of a level one cache memory in response to a level one cache memory write hit and a level two cache memory hit as recited in the claim.

Lai discloses invalidating a write data element of a level one cache memory in response to a level one cache memory write hit and a level two cache memory hit [*hit in L2 cache and hit in L1 cache, corresponding cache line can be invalidated*; col. 2, lines 41-57] to maintain multiprocessor coherency.

Since the technology for implementing invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit was well known and since invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit maintains multiprocessor coherency, an artisan would have been motivated to implement invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit in the system of AAPA and Hazawa. Thus, It would have been obvious to one of ordinary skill in the art, at the time the invention was made, to modify the system of AAPA and Hazawa to include invalidating a corresponding level one cache memory location in response to a level one write hit and a level two write hit

because it was well known to maintain multiprocessor coherency and to maintain cache coherency between the primary cache and main memory (col. 4, lines 7-10) as taught by Lai.

As per claim 24, Lynch discloses a snooping circuit [*snoop queue 402*; Fig. 4].

As per claim 25, Lynch discloses said write data element is located within said level one operand cache memory [col. 3, lines 43-45].

Conclusion

23. The examiner requests, in response to this Office action, any reference(s) known to qualify as prior art under 35 U.S.C. sections 102 or 103 with respect to the invention as defined by the independent and dependent claims. That is, any prior art (including any products for sale) similar to the claimed invention that could reasonably be used in a 102 or 103 rejection. This request does not require applicant to perform a search. This request is not intended to interfere with or go beyond that **required** under 37 C.F.R. 1.56 or 1.105.

The request may be fulfilled by asking the attorney(s) of record handling prosecution and the inventors)/assignee for references qualifying as prior art. A simple statement that the query has been made and no prior art found is sufficient to fulfill the request. Otherwise, the fee and certification requirements of 37 CFR section 1.97 are waived for those documents submitted in reply to this request. This waiver extends only

to those documents within the scope of this request that are included in the application's first complete communication responding to this requirement. Any supplemental replies subsequent to the first communication responding to this request and any information disclosures beyond the scope of this are subject to the fee and certification requirements of 37 CFR section 1.97.

In the event prior art documentation is submitted, a discussion of relevant passages, figs., etc., with respect to the claims is requested. The examiner is looking for specific references to 102/103 prior art that identify independent and dependent claim limitations. Since applicant is most knowledgeable of the present invention and submitted art, his/her discussion of the reference(s) with respect to the instant claims is essential. **A response to this inquiry is greatly appreciated.**

24. The examiner also requests, in response to this Office action, support be shown for language added to any original claims on amendment and any new claims. That is, indicate support for newly added claim language by specifically pointing to page(s) and line no(s) in the specification and/or drawing figure(s). This will assist the examiner in prosecuting the application.

25. When responding to this office action, Applicant is advised to clearly point out the patentable novelty which he or she thinks the claims present, in view of the state of the art disclosed by the references cited or the objections made. He or she must also show how the amendments avoid such references or objections See 37 CFR 1.111(c).


Art Unit: 2188

26. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Pierre M. Vital whose telephone number is (571) 272-4215. The examiner can normally be reached on 8:30 am - 6:00 pm, alternate Fridays off.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Mano Padmanabhan can be reached on (571) 272-4210. The fax phone number for the organization where this application or proceeding is assigned is 703-872-9306.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

May 5, 2005



Pierre M. Vital
Primary Examiner
Art Unit 2188